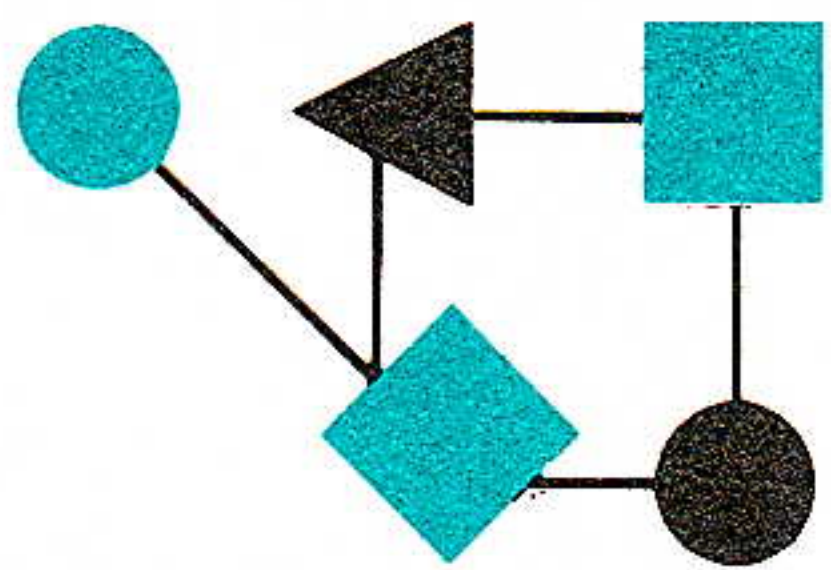


CONNEXIONS[®]



The Interoperability Report

August 1992

Volume 6, No. 8

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

Comparing Compound Document Processing Models..	2
The Guy with the Bike.....	10
Multiprotocol Encapsulation over Frame Relay.....	14
Internet Humor.....	17
KRFC AM 1313.....	18
Comments on: "OSI is Still a Good Idea".....	20
Book Review.....	22
Announcements.....	23

From the Editor

A year ago we brought you an article on ODA, the *Open Document Architecture* as specified in ISO 8613. ISO has also published another standard for the open interchange of revisable documents, namely SGML, the *Standard Generalized Markup Language* (ISO 8879). In this issue, Mark Bramhall and Jon A. Stewart of Digital Equipment Corporation offer a comparison of these two standards, although they note that this is perhaps a bit like comparing apples and oranges.

Carl Malamud's new book *Exploring the Internet: A Technical Travelogue*, is now at the printer and will be ready for distribution at INTEROP 92 Fall in October. This month we have another extract from the book, as Carl visits Steve Roberts and "The Bike" at Nomadic Research Labs in Mountain View, California.

Frame Relay has rapidly emerged as an efficient method to interconnect communications devices such as routers and bridges over the wide area. The need to ensure interoperability over Frame Relay between LAN internetworking devices has motivated the development of a common encapsulation protocol. Work on this protocol has been performed under the auspices of the IETF, and resulted in RFC 1294, "Multiprotocol Interconnect over Frame Relay." Andrew Malis from BBN Communications gives an overview of this technique, starting on page 14.

If you've been reading messages on the IETF mailing list recently, you've undoubtedly discovered that the Internet community is knee-high in politics. Messages (or should I say "flames"?) clearly bare witness to the fact that internetworking has become *very serious* business. I thought it might be appropriate, therefore, to remind everyone about the brighter side of this community, and asked Garrett Wollman to write a piece about Internet humor. As an example of the true Internet spirit, we've also included RFC 1313.

Marshall Rose responds to Dick desJardins' article "*Opinion: OSI is (Still) a Good Idea*" which appeared in our June issue. The OSI debate does not seem to have died down completely yet, and I have promises for more articles from other quarters. Stay tuned.

I did mention, in the June issue, that a *Special Issue* on Electronic Mail and Directory Services was planned for publication this month. Well, it's almost ready to go, and I promise it will be the next issue (September).

ConneXions is published monthly by Interop Company, 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP. E-mail: connexions@interop.com.

Copyright © 1992 by Interop Company. Quotation with attribution encouraged. ConneXions—The Interoperability Report and the ConneXions logo are registered trademarks of Interop Company.

Comparing Compound Document Processing Models

by Mark Bramhall and Jon A. Stewart,
Digital Equipment Corporation

Introduction

The International Organization for Standardization has so far published two standards for the open interchange of revisable documents. The first was *SGML*, the *Standard Generalized Markup Language* (ISO 8879). The second was *ODA*, the *Open Document Architecture* (ISO 8613). Vendor-based models, such as Digital Equipment Corporation's *Compound Document Architecture* (CDA), as well as a range of proprietary products, including word processors and desktop publishing systems, enjoy *de facto* recognition as standards from a broad base of users and third party developers. It's been said that comparing these standards is like comparing apples and oranges. Not only do they do different things, but often they do the same things differently.

Ask an expert whether two standards offer the same feature and the most likely response is: "It depends." Ask for a checklist of features on which to base a comparison, and the number of qualifying questions may well exceed the number of features: Do you include references to external standards as part of your model? Do you need to maintain a consistent structure throughout all stages of document processing? Would a standard that lacks a defined process for document revision qualify as a standard for interchanging revisable documents?

These qualifications can contribute more confusion than precision. The details of how models work, and therefore how they differ, is already available. What users need is a clear overview of significant model characteristics and a consistent basis of comparison. Even the word "model" used here is a struggle to find a word that does not preclude this or that member from what is decidedly a mixed bag of architectures, languages, word processors, and more. It is ironic that a topic so fundamental to the basic interchange of information in the computer age should itself be so difficult to talk about.

One way to find a general method of comparing apples and oranges is to go through the process of actually comparing an apple and an orange to see if there aren't general rules that seem to consistently distinguish different kinds of fruit. What follows here then is: 1) an examination of SGML and ODA, 2) a recommended approach for making general comparisons between models, and 3) an application of the approach to a comparison of four ways to interchange revisable documents: SGML, ODA, Digital's CDA, and a fourth way called "private formats," used by proprietary products.

SGML and ODA compared

Key differences between SGML and ODA show up right in their names. SGML is a "language" for defining markup directives. ODA is an "architecture." What's the difference? A language is a set of symbols (a vocabulary) and a set of rules (a grammar) for using those symbols to convey meaning. An architecture is a plan for an integrated system of components that carry out some purpose. In other words, an architecture can do other things besides convey meaning. It could, for example, provide a way to revise documents, not simply encode documents that have already been revised.

It is often easier and faster for vendors to implement a language than it is an architecture. (Which may explain why SGML is widely implemented today, while there are still very few ODA products.) It is easier to design a language that supports "all" applications than it is to actually design "all" applications. An architecture, however, can bring more solutions to bear on different parts of a problem, and with a greater likelihood that these various solutions will work together.

Used as a markup language, SGML presumes documents are character-based text. No other data types are specifically accounted for in the standard. Historically, the term “markup” refers to the extra symbols that a copy editor writes on a galley to tell a typographer how to lay copy out as a document. The concept of markup is therefore rooted in text, and in particular the layout of text. In its SGML usage, markup is employed to guide a computer program that is performing some process on document content. SGML itself does not specify what to do with the marked up content.

ODIF One of the most apparent indications of SGML’s textual orientation is the fact that an SGML encoding is itself character text that can be read and worked on directly by human beings. Users have even been known to apply SGML encoding themselves rather than, say, rely on an automated word processor. The only ODA encoding used in current products is computer oriented binary code. This encoding is called ODIF (*Open Document Interchange Format*.) It is not human readable and must be processed by an application to be turned into its readable (and final) forms. (A text encoding of ODA-compliant documents is also specified in the standard. This encoding is based on an application of SGML called *Open Document Language*.)

ODA The ODA standard defines a document as some amount of structured information “intended for human perception that can be interchanged as a unit between users and/or systems.” There is no specific text orientation either in terms of document encoding or content. ODA uses ASN.1 (ISO 8824:1987, Specification of Abstract Syntax Notation One, and ISO 8825:1987, Specification of Basic Encoding Rules for ASN.1) to describe documents as instances of objects and to specify their content, logical relationships, physical layout, and presentation style. (Ed.: See how this is done in detail in [1] and [2]).

Logical and Layout structures

An ODA/ODIF document encoding results from ASN.1 object class, object attributes, object relationships, and object content definitions. Despite the fact that it is quite general in terms of what it will let you do to documents, ODA is very specific in terms of allowable object classes, attributes, relationships, and methods. For example, permissible content types include ISO character sets, geometric graphics (based on CGM, the ISO standard for computer graphics meta files), and raster graphics (based on CCITT’s facsimile recommendations). ODA defines two kinds of objects: logical objects and layout objects. ODA documents, therefore, can have both a logical structure and/or a layout structure. The fact that a document can have separate logical and layout structures—each different but well-defined—lets a document be reformatted (by overriding the layout structure) while maintaining as much of the author’s original intent as possible.

The relationships among logical objects (e.g., chapter, paragraph) make up a document’s logical structure. The relationships among layout objects (e.g., pages, sidebars) form the document’s layout structure. A generic logical structure describes the rules for relationships among logical-content objects, such as sections and chapters (e.g., a section can be part of a chapter but not vice versa). The specific logical structure is nothing more than a particular instance of a document that obeys these rules (e.g., this particular chapter contains these particular sections). Similarly, generic layout structure is a set of rules that control layout, and a specific layout structure is also a document instance—it is the end result of a specific logical structure that has gone through the layout process for presentation into pages, text blocks, numbered sections, etc.

Document Processing Models (*continued*)

A document's logical structure is always preserved, even in the face of changed layout, because the logical structure contains the content—the words, figures, etc.—created by the author. It is often very useful to maintain the logical structure while changing the layout structure, such as when you want to re-format a document for presentation on different devices (character versus *PostScript* printer, or printed manuals versus on-line documentation). So layout objects, such as pages, can be described as the way to provide a presentation context within which to present logical objects.

Both logical and layout objects have attributes that guide the layout and presentation of documents. These attributes are 1) *presentation attributes*, and 2) *layout directives*. Only logical objects that contain content, however, can provide presentation attributes. An example of a presentation attribute might be *italics* or a specific font. Layout directives are the means for a logical object to control its own layout. An example of a layout directive might be that you can't break figures (a logical object class) across pages (a layout object class). All layout/presentation attributes are typically collected into style "objects" that are referenced by the object identifiers of the objects that use them. The freedom to reference (and therefore send, receive, revise, apply) layout independently from content, is characteristic of a document architecture.

SGML

In its usage within SGML, markup does not necessarily have to be applied only to text—so long as the application knows how to handle the content type referred to by the markup. Nor does the application necessarily have to be layout or text processing. An SGML-capable application may be a document formatter—an application that "hides" the markup symbols and positions the content for display or printing. However, the application might have nothing at all to do with document formatting, as in the case of a content retrieval system. An important difference between SGML and architectures such as ODA is that SGML does not define a document standard—either in terms of allowable content (the content model) or what can be done to that content (the process model).

Markup can indicate the placement of other data types, such as raster images, graphics, and even sounds, within a document. However it is up to some application to provide a process for handling this other content, for binding the content into the main document, and for integrating the processes supporting the various content types. For example, an SGML text editor may "escape" to a TIFF graphics editor when the user selects a graph that is to be edited within the copy. Later, the TIFF process must allow the user to escape back to the main document to complete editing. (This return is one of the things not defined by SGML and leaves the door open for incompatibilities between SGML applications.) Markup can also be used to indicate other operations on content besides layout. One example of a non-layout markup function would be to mark some string (that is, an information object such as title, author, etc.) for future retrieval from a database.

GIs

SGML is a language in the sense that it employs both a vocabulary of markup "tags" and rules of "grammar" that determine how these tags can be combined to create document structures. The tags are called *generic identifiers*, or GIs, and are written directly into the clear text of a document. Typically, GIs are written within angle brackets, such as <para> to indicate a paragraph or <author> to indicate an author.

An application uses the GIs to determine what processing to apply against the content (appearing in line) referenced by the GIs. Non-text entities are also indicated by GIs.

DTDs

The grammar of SGML is the rules for defining the representation (e.g., as integers, ASCII characters, real numbers, etc.) and use of GIs (e.g., a chapter heading can only occur at the beginning of a chapter). In ODA, these rules are defined directly in ODIF as allowable ASN.1 syntax. SGML, however, does not itself define the meaning of either its tags or grammar. Those definitions are provided by external specifications, called *document type definitions* (DTDs).

In order to process a document, an application must understand a document's markup. This means both the document and the application must abide by the same DTD, as must any group of users who intend to interchange or revise each other's SGML encoded documents. It is entirely possible that two users will both employ SGML and be completely incompatible because they use different DTDs. To achieve "standard" document interchange and revisability, groups of users adopt common public DTD libraries. There are hundreds of such libraries—which can mean that there are hundreds of "standard" SGML implementations—none of which are compatible! Well known DTD libraries include:

- CALS: (*Computer-aided Acquisition and Logistics Support*), a DoD specification (MIL-M-28001)
- AAP: *American Association of Publishers*
- AIA: *Aviation Industry Association*
- AMA: *American Mathematical Association*

In summary, SGML is a language with a very low-level vocabulary and grammar which are used to define higher-level markup languages. The "L" in SGML is defined in the DTDs. SGML only standardizes what is a valid DTD, and only DTDs standardize what is valid markup for a specific text document. By contrast, ODA does not have to go "outside" for semantics with which to encode or interchange documents—including documents with non-text data. ODA does not, however, provide its own method for processing those semantics.

Architecture of Architectures

By relying on external semantics, SGML gives up compatibility between document processing applications in exchange for the ability to define these applications. In some sense, SGML is an "architecture of architectures." One kind of universality is achieved in exchange for another. By providing its own semantics, ODA deliberately limits the behavior of applications for the sake of compatibility. But not specifying a process model to implement those semantics means ODA can not always guarantee that compatibility is achieved. One conclusion that can be drawn from this is that going outside or staying inside for either semantics or process makes a model better for some things, worse for others. Perhaps SGML and ODA could each do what the other does if SGML also defined its own semantics and process model and ODA also allowed for external semantic definitions.

A model comparison

Comparing SGML and ODA is actually worse than comparing apples and oranges. At least with apples and oranges you can easily reach a common level of abstraction (fruit) and agree on common attributes of comparison, such as weight, vitamin content, and calories. Attempting to apply any generality to SGML and ODA, on the other hand, can cause endless qualifications—to the point where a common standard of comparison can create more confusion than it eliminates.

continued on next page

Dimensions

These differences fall on two dimensions: outside/inside and semantics/process. On the vertical axis, the outside/inside dimension indicates whether the capability is addressed within the model (cells A and D) and/or with external references (cells B and C). The semantics/process dimension is on the horizontal axis. A model is more complete when it allows the user to convey intent and provides its own process for carrying out intent. It is more inclusive if it allows the user to invoke alternative external means of conveying intent and alternative external means for carrying out the intent. In other words, the stronger the model the more cells that are shaded.

These cells cover four possibilities:

- That the model's own semantics are sufficient to convey the users intent
- That the user has the option of using alternative semantics (say, as defined by a DTD)
- That the model lets you select from alternate processes (such as to format data)
- That the model provides its own process

The selection of these two particular dimensions is an attempt to find consistent and general criteria that distinguish among models that legitimately claim to do the same thing but do it differently. It is also an attempt to settle arguments over which way is better. For example, arguments can be made that a model (like SGML) that supports alternative semantics to revise documents is stronger because it is less confining. Another argument might say that the model is weaker because the model leaves open the possibility of incompatible implementations. Figure 1 would say that supporting external semantics (or processes) and providing your own are not mutually exclusive. In fact, a stronger model would do both. If the user can't have both then the matrix at least tells you what's missing to enable comparison based on individual need.

The matrix provides a snapshot comparison of both relative model strength and capability importance. A stronger model has fewer "missing pieces" or unshaded cells. A more important capability is one that has broader support from more of the models. By rank ordering the models and the benchmarks so that the most white space occurs toward the right and downward, the most capable models end up being listed toward the top and the most important capabilities are listed to the left. (Actually the only capability which demonstrated less importance—as measured by the breadth of support it received from the models—was the ability to separate content from layout. However, if the category, *Private Formats*, were removed from the comparison, this capability would move from last to first place in importance.)

Choice of capabilities

The choice of capabilities is open ended and not necessarily limited to those listed here. The better the selection of capabilities chosen for comparison, however, the more revealing the comparison among models. More can be learned, for example, by comparing what might be called general capabilities than by comparing derived capabilities. Document interchange is an example of a general capability. Blind interchange is an example of a derived capability. Blind interchange is possible if a model provides its own interchange semantics rather than only relying on external semantic sources.

Document Processing Models (*continued*)

Only a comparison on the general capability of interchange reveals this referencing difference among models. A comparison on the derived capability of blind interchange support reveals only that one model does and one doesn't—with no clues as to why.

Observations

Only a few general capabilities need be selected in order to provide a clear comparison among models or an overview of their characteristic approaches to document processing. Although a 4 x 4 model/capability matrix provides 64 separate cells for comparison (and potential discussion), the model focuses attention on a select number of general observations.

Among these observations:

- That CDA is the most “complete” architecture on these criteria, providing both semantic and process support for each capability listed, while giving the user the flexibility to go outside the model for special features or to achieve compatibility with other models.
- That the Private Formats approach tends to be least inclusive. Proprietary products rely entirely on their own semantics and process for features. If any interchange or cooperation happens, it results from homogeneous applications and platforms using a pair-wise agreement on what the data format is. Also there is limited ability to handle layout separately from content.
- That ODA is more complete (on the criteria listed) than SGML, having almost twice as much “capability coverage.”
- That ODA and SGML are chiefly distinguished by the latter's lack of semantics with which to encode the user's intent.
- That capabilities distinguish models, not vice versa. (Each model usually displays the same pattern of white and shaded cells for each capability—i.e., a model will generally employ a characteristic support strategy across capabilities.)

Conclusion

No universal rule book exists for the creation or selection of compound document processing models. Developers, sponsors, advocates, and users are each guided by their own objectives. These can be political, economic, and technical. Each model, therefore, must be evaluated in terms of the individual agenda of the person doing the evaluating. Whether a particular tool is a language or an architecture is not as important, for example, as whether the tool does what needs to be done.

A close comparison of SGML and ODA reveals that document processing models cannot always be distinguished by whether or not they support various capabilities. Two models may support the same capability and still be very different in the way they provide that support. In general, models are more consistently distinguished by how they support various capabilities. A support strategy can be distinguished by whether semantics and process are provided inside and/or outside the model.

The “signature” characteristics of a model are usually more obvious when looking at general capabilities than at derived capabilities. The existence of a derived capability is often codependent on support of a particular general capability and a particular support strategy. It is therefore unlikely that a derived capability would be a good indicator of support strategies.

It might be left unclear, for example, whether a model lacks a particular feature because it simply wasn't implemented or because its choice of strategies makes implementation difficult.

Applying this method of comparison to four very different kinds of processing models shows that not only is there a way to compare "apples and oranges" but that it is also possible to rank order models according to how many generic "pieces" are missing from both kinds of fruit. Whether a particular piece is missing, however, may not be the critical issue. Knowing how models compare in their overall approach to the general tasks of compound document processing may be more important.

References

- [1] Stewart, J. "Components of OSI: Document Interchange Using ODA," *ConneXions*, Volume 5, No. 8, August 1991.
- [2] Howard Motteler and Deepinder Sidhu, "Components of OSI: Abstract Syntax Notation One," *ConneXions*, Volume 6, No. 1, January 1992.
- [3] International Organization for Standardization, "International Standard ISO 8879—Information processing—Text and office systems—Standard Generalized Markup Language (SGML)," 1986.
- [4] International Organization for Standardization, "International Standard ISO 8613—Information processing—Text and office systems—Office Document Architecture (ODA) and interchange format," 1989.

Part 1: Introduction and general principles

Part 2: Document structures

Part 3: (no longer exists)

Part 4: Document profile

Part 5: Office Document Interchange Format (ODIF)

Part 6: Character content architectures

Part 7: Raster graphics content architectures

Part 8: Geometric graphics content architectures

[Note that the cited documents say "Office Document Architecture" rather than "Open Document Architecture." The name change to "Open" has been approved, but the new documents will not be published until this summer.]

MARK H. BRAMHALL is a Consulting Software Engineer at Digital Equipment Corporation. Having joined Digital in 1970, he has worked with all of Digital's major operating systems and processors. He currently directs the CDA architecture group, focusing on multimedia, compound document solutions, integration, and futures. Previously, he worked on text editors, human factors research, and operating systems for the PDP-8, PDP-11, and VAX. He has participated in X/Open and ECMA and represented Digital in the negotiations and forming of the ODA Consortium. He can be reached as: bramhall@ddif.enet.dec.com.

JON STEWART received his BS in Physics (Tulane University, 1956) and MA in Astronomy/Astrophysics (Indiana University, 1958). He programmed stellar model computations at IU as a graduate student and was also involved at IU in some of the earliest (1961) research in computational linguistics. He has been a computer science educator and research associate/consultant, a DEC software support and educational marketing specialist; and, for DEC since 1978, a software designer. Since 1985 he has been involved in establishing DEC's international software engineering architectures and standards, analyzing potential international software alliances, and representing DEC in the development of international standards committees for document interchange and processing. He is the Chair of the task group in ANSI X3V1 that coordinates the USA's participation in ODA. He can be reached as: stewart@ddif.enet.dec.com.

The Guy With the Bike

by Carl Malamud

[Ed.: In this issue, Carl Malamud the *Internet Explorer* goes to Mountain View to talk to "The Guy With the Bike," part of the three trips around the world Malamud is taking to research his next book. The result of all this travel will be *Exploring the Internet: A Technical Travelogue*, published by Prentice Hall and distributed to all conference attendees at INTEROP 92 Fall in San Francisco].

Saturday afternoon, I headed down to Mountain View, the home of Sun Microsystems. I was going to visit Steve Roberts, better known as "The Guy With the Bike."

On the road

In 1983, Steve decided that a life of consulting and writing books on the subject of microprocessors was not for him. He tallied up the things he liked to do, and decided on a short list of writing, meeting people, and riding bikes. He sold his home in Columbus, Ohio and invested everything he had in a recumbent bicycle, an early laptop, solar panels to power his ham radio and other devices, and then hit the road.

For 10,000 miles, he kept on pedaling, going through Florida, the Rockies, the deserts of Utah and California, and the bleak desolation of San Clemente. On the way, he used *CompuServe* as his electronic mailbox and wrote articles for any publication that would send him money. The bike was certainly a natural conversation opener, helping to fulfil his most important ambition of meeting lots of people.

The original Winnebiko was only a couple of hundred pounds and had 18 gears. Over time, that original frame has grown to support 105 gears, 580 pounds, and an incredibly sophisticated collection of on-board computers.

BEHEMOTH

Sun Labs has loaned space to Steve to let him work on the latest incarnation of the Bike, known as "BEHEMOTH," the *Big Electronic Human-Energized Machine...Only Too Heavy*. I reported to Sun's building 4, and was met by Steve, a tall bearded man in his late thirties, wearing a "Peace" T-shirt. He led me through dark empty corridors to a locked door bearing a neatly labeled sign which read "Nomadic Research Labs."

"Labs" is an understatement in this case. Under one of the desks is a futon and scattered throughout are signs that the lab is also a home. Walking in the first time, though, you don't notice the futon, the flute, or the cereal bowl. Your eye naturally turns to the huge bike in the middle of the floor.

At first glance, you see a recumbent bicycle with a storage unit mounted behind the seat and an aerodynamic hood up front. Behind the bike is a trailer, covered with 72 watts of solar panels, used to charge 45 amp-hours worth of batteries.

Communications

Sticking up from the rear of the trailer is a big yellow pole which contains the ham radio antennas. The pole is six feet tall and can be extended to 12 feet for even better reception. Another antenna on the back is a little Qualcomm satellite dish, similar to what you would use on a delivery truck. Roberts keeps a Sun 4/260 at Qualcomm headquarters and runs 165 bps *uucp*-based transfers into the bike.

The satellite dish provides coverage anywhere, but 165 bps isn't really optimal for true connectivity. A high-speed modem and a cellular phone allow high-speed IP access to a server located at Sun.

Computers

On the bike, there are a variety of different computers. A Sparcstation, a Mac portable, and a DOS portable have all been ripped apart and mounted into the bike. The screens are mounted on the bike console, with one screen flipping up to expose another. (Steve likes to say that this is "Mechanical Display Paging").

Another laptop is in a carrying case that can be removed, allowing Steve to compute in diners and other places that the bike is not welcome. The carrying case is, of course, in constant communication with the bike.

Input

Input to computers while riding is provided by a handlebar keyboard and a head mouse. The handlebar keyboard is based on an Infogrip BAT chord keyboard which has a total of 7 keys. Force sensing resistors are built into the handlebar grips and are linked to the BAT keyboard controller.

The chord keyboard allows Steve to type at roughly 35 words per minute, not an optimal rate when you make your living being paid by the word. Once the keyboard controller generates characters, the data is sent to a macro package on a DOS machine which expands the data. Through careful use of macros, the data hitting the target computer appears to be a typist working at 100 words per minute.



Steve Roberts with BEHEMOTH

The head mouse is the other major form of input. Three transducers are mounted on Steve's helmet and they are used to sense an ultrasonic beam generated from the console. The resulting motion that is detected is fed into the controller of a Macintosh ADB mouse.

Output

The helmet has a few other features to help the digital nomad. A 720 x 280 pixel red image floats in front of him on a little display mounted on the helmet and acts as a console. A heat exchange system built into an inner liner for the helmet acts as an air conditioner, using the refrigeration unit on the bike to cool fluids which travel up and remove heat from the helmet.

continued on next page

The Guy With the Bike (*continued*)

The helmet also has a little microphone on it which connects to a serial cross-connect bus. The bus can set up connections to cellular phones, the ham radio, or even the voice recognition unit on the bike. The cross-connect bus also links modems to serial ports and any other serial links that need to be made. The audio bus can handle up to 8 simultaneous events and the serial bus can handle up to 4 simultaneous events.

Power

Just as the serial bus goes throughout the bike, so does a power bus. Power is distributed in a series of batteries. The main power is three 15 amp-hour batteries, plus there are various other special purpose batteries scattered about.

On sunny days, photovoltaic cells recharge the system. Additional power is provided by a regenerative braking controller, which transforms the heat generated by braking into power. If there are no hills and no sun, a cable connected to a car's cigarette lighter will do the trick.

A Motorola 68HC11, programmed in *Forth*, acts as the power controller, sending power where needed and monitoring usage. The controller will signal the main bicycle control unit when power is low, which in turn tells Steve so he can take corrective action. Two other 68HC11 controllers handle the serial and audio buses.

The audio bus allows sound output from many sources to be mixed. Two 4" Blaupunkt speakers are mounted behind the driver and connect to an automotive stereo amp. Data from the CD player, the cellular phone, the ham radio, or the speech synthesizer can all be mixed. Steve carries over 100 CDs with him on the road, removed from their original packaging to save space and weight.

Security

This entire contraption is now worth over \$300,000 in parts alone. If you add time, it is easily worth over \$1 million. Security naturally becomes an issue.

Security for the BEHEMOTH is provided on several levels. At the most basic level, a microwave doppler motion detector reports any motion within 8 to 10 feet. Other detectors signal when a person touches the bike or sits on the seat. A *Global Positioning System* (GPS) receiver on the bike hood reports bike movement and speed.

Different responses can be set to security events depending on the circumstances. The voice synthesizer can use the speakers to utter an appropriate phrase if somebody touches the bike, such as "do not touch, or you will be vaporized by a laser beam!" A siren can be activated, or Steve can be paged.

If the bike starts moving without a password, more drastic actions are possible. The wheel can lock, a call can be placed to 911 and the speech synthesizer can cry for help, the siren can go off, and, most effective of all, a wild-eyed digital nomad is automatically programmed to come bursting out of his tent in a frenzy.

Attention

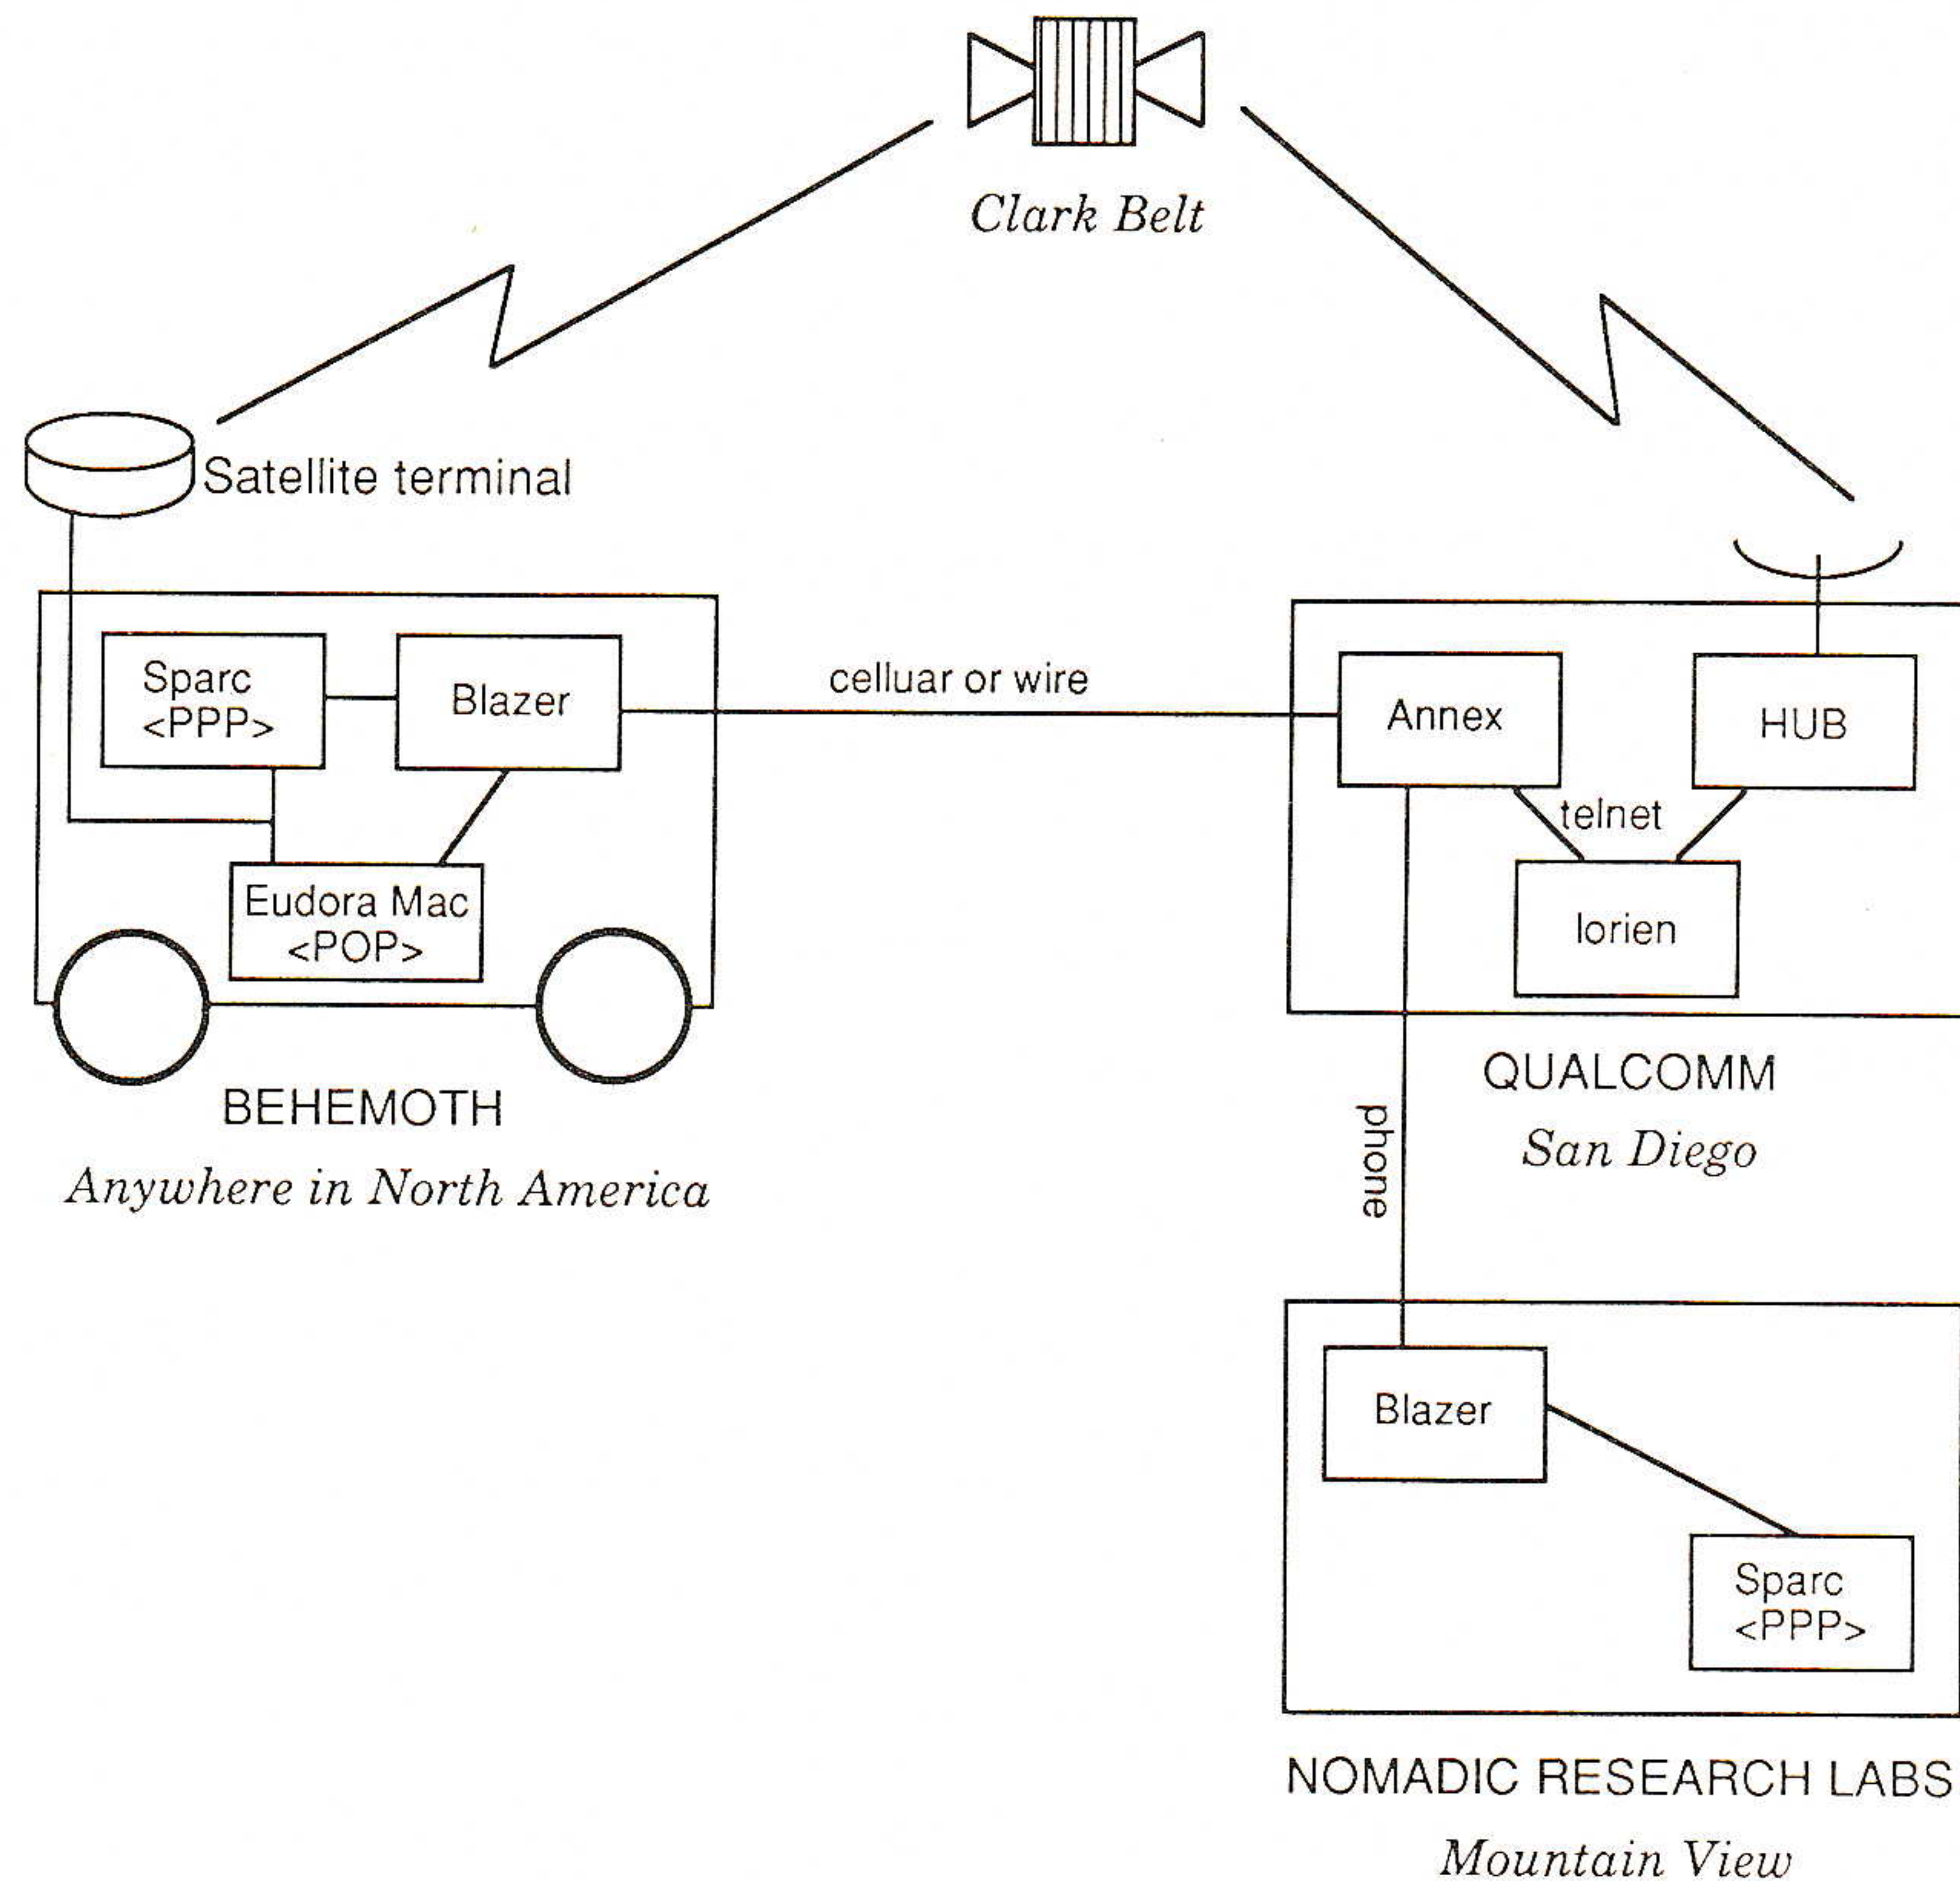
The BEHEMOTH and its previous incarnations have attracted the media from the word go. CBS, *USA Today*, and even the trade press quickly realized that this was a story. Being "The Guy With the Bike" has made Steve instantly recognizable throughout the US.

For the first 17,000 miles of journeys, Steve simply wandered. An office manager at home base took care of things in return for a cut of the revenue.

The Mothership

Over time, though, Nomadic Research Labs has grown into quite a business. Steve publishes a “sometimes quarterly journal” called *No-madness* and is in frequent demand on the lecture circuit, providing an interesting alternative for groups that can’t afford Oliver North.

To meet the demands of speaking engagements and trade shows, the bike now has a mothership, a large trailer pulled by a nice van. The bike goes inside the mother ship and hooks up to the antennas, and a little console goes up in the front seat. The bike continues to be the communication hub, even though it is inside a trailer.



BEHEMOTH communication system

I thought that my three round-the-world trips for *Exploring the Internet* were quite a journey, but I felt like a digital dilettante compared to Steve’s wholehearted commitment. His lifestyle is different from that of your usual commute-to-cubicle engineer, but he has certainly proved that you can have fun and put together a technical tour-de-force at the same time.

I left Steve to drive up the peninsula to San Francisco to meet Ole Jacobsen, my editor at *ConneXions*. Over a dinner of pasta with garlic, olive oil, and anchovies, we plotted how we could get Steve Roberts to pedal around the INTEROP show floor, hooking, of course, into the Shownet. How about remote SNMP bike management? A BEHEMOTH MIB? The possibilities were endless.

CARL MALAMUD (carl@malamud.com) has recently finished his world travels (at least for the time being) and settled down in Washington, DC. He is the author of several technical reference books, and enjoys many fine lunches and dinners.

Ed.: The bike is constantly evolving as new technology becomes available and is road-tested. Thus, any detailed description of the system is bound to be out of date by the time you read this. Next month: A review of Steve Roberts’ book “Computing Across America.”

Multiprotocol Encapsulation over Frame Relay

by Andrew G. Malis, BBN Communications

Introduction

Frame Relay has rapidly emerged as an efficient method to interconnect communications devices such as routers and bridges over the wide area. The adoption of *Frame Relay* standards allows the potential for establishing multivendor network solutions. Unfortunately, even if two devices correctly implement the standards necessary to connect to a Frame Relay network, they may not be able to pass data across a network. Although Frame Relay standards specify a standard data link layer protocol, devices built by different vendors must additionally share a common network-layer protocol to effectively interoperate across a Frame Relay network.

In Local Area Network (LAN) environments this is well specified. For example, all routers use a common method of carrying IP packets (a network-layer protocol) over an Ethernet (a data-link layer protocol), which guarantees their ability to interoperate over Ethernet networks. In a Frame Relay environment, however, two routers may both route IP packets, but utilize different procedures for transporting or encapsulating IP packets over a Frame Relay network. In this case the routers are unable to interoperate across the network. Similarly, two bridges may both support identical remote bridging procedures, but if they do not use a common encapsulation format for sending the bridged frames across a Frame Relay network they cannot interoperate.

The need to ensure interoperability over Frame Relay between LAN internetworking devices motivated the development of a common encapsulation protocol. This protocol may also be used in host computers that directly connect to Frame Relay networks. Work on this encapsulation protocol has been performed under the auspices of the *Internet Engineering Task Force* (IETF), and resulted in Internet Request for Comments (RFC) 1294, "Multiprotocol Interconnect over Frame Relay" [1]. RFC 1294 is a Proposed Internet Standard.

Encapsulation description

Since Frame Relay is a wide-area networking protocol, and conforms with ANSI and CCITT protocol standards, RFC 1294 is based upon a network-layer protocol encapsulation that also conforms with existing ANSI, CCITT, and ISO standards. This encapsulation method is specified in ISO/IEC Technical Report (TR) 9577, "Protocol identification in the network layer" [2], and is jointly administered by ISO and CCITT. Because TR 9577 was written mostly for use with X.25, the packet formats in the Technical Report had to be adapted for use with Frame Relay.

Frame Relay frames using this encapsulation method have the following format (shown as a series of octets):

Flag (0x7E)
Data Link
Connection Identifier (DLCI)
Control (0x03)
Optional Pad octet(s) (0x00)
Network Layer Protocol ID (NLPID)
•
•
Data
•
•
Frame Check
Sequence (FCS)
Flag (0x7E)

In the diagram, numbers are shown as hexadecimal values. Each frame, as usual, begins and ends with a flag octet. The first two octets following the beginning flag contain the normal Frame Relay addressing and control information (DLCI, BECN, FECN, etc.).

The next octet, labeled "Control," is not strictly required for Frame Relay networks, but is included for complete compatibility with CCITT standard Q.922 [3], of which Frame Relay is a subset. As used in this encapsulation, the Control field is set to 0x03, which identifies the frame as an "Unnumbered Information" frame.

The next field is optional. The frame may contain 0 or more padding octets, each of which is set to 0. This padding optionally aligns the remainder of the frame to a convenient memory boundary for the sender (usually to optimize the sender's performance).

The next field contains the *Network Layer Protocol Identifier* (NLPID). This field is critical to interoperability over Frame Relay because it allows the identification of network layer protocol types between devices on a network. The NLPID is specified by TR 9577, and can contain the following hexadecimal values:

0xCC: Internet Protocol (IP) [4]

0x81: ISO Connectionless-mode Network Protocol (CLNP) [5]

0x80: IEEE Subnetwork Access Protocol (SNAP) [6]

There are other legal values defined, but these are the ones most commonly used for LAN interconnection. The value 0 is illegal, so that the NLPID field can be distinguished from the preceding optional padding.

The use of IEEE SNAP allows a large number of other protocols to be encapsulated. SNAP defines a five-octet header, and is used on LANs such as Ethernet, Token Ring, and FDDI for protocol identification. This allows LAN packets to be encapsulated almost verbatim to be sent over Frame Relay.

SNAP also contains IEEE-defined values to allow many kinds of remote bridging. These values are listed in RFC 1294, along with the frame formats used bridging various media types.

The next field following the NLPID octet is the *Data* field. For IP and CLNP, this is simply the encapsulated datagram. When using SNAP, this is the five-octet SNAP header, followed by the encapsulated LAN packet or bridged frame. RFC 1294 describes frame formats for each of the various encapsulations.

Following the data field are the normal Frame Relay *Frame Check Sequence* (FCS) and closing flag octets.

Besides defining the standard encapsulation for network layer protocols over Frame Relay, RFC 1294 also sets standards for data link layer parameter negotiation, packet fragmentation and reassembly, and network layer address resolution over Frame Relay.

Current status

RFC 1294 was reviewed during its development by the IETF, the Frame Relay Forum Technical Committee, and the ANSI committee responsible for TR 9577. The RFC is now a Proposed Internet Standard, and has been implemented by a number of router and bridge vendors [7]. Several of these implementations have already been tested together to confirm interoperability.

Encapsulation over Frame Relay (*continued*)

Future plans

RFC 1294 will follow the normal cycle for becoming an Internet standard [8]. A *Proposed Internet Standard*, such as RFC 1294, can be advanced to the status of *Draft Internet Standard* after at least six months have passed, several interoperable implementations have been demonstrated, and operational experience has been obtained. A Draft Internet Standard can be advanced to the status of full Internet Standard after at least four additional months have passed and significant implementation and operational experience have been obtained. Comments and suggestions for improvements to the standard are solicited at each step in the process.

Besides following the Internet standards track, the RFC has been converted to a draft ANSI standard (Annex F to ANSI T1.617 [9]) and, once approved by ANSI, will be submitted to ISO and CCITT for their approval.

To obtain RFC 1294, RFCs are available in a number of public electronic archives and by electronic mail. For further information on obtaining RFCs, send Internet electronic mail to `rfc-info@isi.edu`, with the message body "help: ways_to_get_rfc".

References

- [1] Bradley, T., Brown, C., & Malis, A., "Multiprotocol Interconnect over Frame Relay," RFC 1294, January 1992.
- [2] Information technology—Telecommunications and information exchange between systems—Protocol identification in the network layer, ISO/IEC TR 9577: 1990 (E) 1990-10-15.
- [3] International Telegraph and Telephone Consultative Committee, "ISDN Data Link Layer Specification for Frame Mode Bearer Services," CCITT Recommendation Q.922, April 1991.
- [4] Postel, J., ed., "Internet Protocol," RFC 791, September 1981.
- [5] "Information processing systems—Data communications—Protocol for providing the connectionless-mode network service," ISO/IEC 8473, 1988.
- [6] IEEE, "IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture," IEEE Standard 802, 1990.
- [7] "Special Report: Frame Relay Products," *Data Communications*, May 1992, pp. 65–86.
- [8] Chapin, L., ed., "The Internet Standards Process," RFC 1310, March 1992.
- [9] American National Standards Institute, Telecommunications—Integrated Services Digital Network (ISDN)—Digital Subscriber Signaling System No. 1 (DSS1)—Signaling Specification for Frame Relay Bearer Service, ANSI Standard T1.617, 1991.
- [10] Kozel, E., "The Cisco/DEC/NTI/StrataCom Frame Relay Specification," *ConneXions*, Volume 5, No. 3, March 1991.

Author's address:

Andrew G. Malis
BBN Communications
150 Cambridge Park Dr.
Cambridge, MA 02140
USA
Tel.: 617-873-3419
Fax: 617-873-4086

ANDREW MALIS received the B.S. and M.S. degrees in computer science at Brown and Harvard, respectively. He joined Bolt Beranek and Newman Inc. in 1979, working on the original ARPANET IMP and (among other things) enforcing the transition from NCP to TCP in the ARPANET switches. He is currently a Division Engineer at the BBN Communications Division, where he is responsible for coordinating BBN's use of Frame Relay and other advanced data services in its products. He also represents BBN at the Frame Relay and ATM Forums and at the Internet Engineering Task Force, and actively participates in the Internet standards process. His Internet mailbox is: `malis@bbn.com`.

Internet Humor

by Garrett Wollman, University of Vermont

Introduction

Some time ago, I was discussing with some friends what the difference was between the Internet standards and the standards promulgated by former standards bodies. It wasn't so much the technical content as the fact that there are people out defending the Internet standards (say, for example, RFC 822, the e-mail message format standard) with a religious fervor that doesn't seem to extend to many ISO standards. Why should this be, I wanted to know. I thought about it for quite a while, and came up with a few reasons.

Openness

The Internet standards process is *open* in a manner that goes far beyond what most marketing people mean when they use that word. Most of the standard-making activities in the Internet take place through a formal process, starting out in the IETF (*Internet Engineering Task Force*). How does one become a member of the IETF? There is no formal membership; you simply join the mailing list and participate in the activities of a working group, and maybe even go to the quarterly physical IETF meetings. Anyone with the technical expertise to make a contribution can, which makes the Internet process vastly more open than most others (where you usually have to be a member of some other organization, that must choose you as its representative).

Community

Thanks to this organization, and to the presence of many interesting and approachable people making visible contributions to the process, the Internet carries with it a sense of *community*, which is lost in the cold formality of *de jure* standardization. And not only that, but Internet standards are actually field-proven *before* they are promulgated, rather than putting out a standard and then revising it later because it was found to be unworkable. Who is the Dave Mills of OSI Time Service? Does OSI even *have* a network time service? If not, how would I go about proposing one?

The Internet standards process seems to be the only one which can laugh at itself in public, and with some regularity. Where else do you see April Fools' standards documents published? Even when the standards are serious they often include a bit of humor to keep readers from being bored to tears. Here are some of the more interesting ones:

Annotated Bibliography

- [1] R. Merryman, "ARPAWOCKY," RFC 527, 22 June 1973. A spoof of Lewis Carroll's "Jabberwocky" with much of the flavor of the early ARPANET.
- [2] Mark R. Crispin, "Telnet Randomly-Lose Option," RFC 748, 1 April 1978. *You too can permit or refuse to permit your computer to crash at any time.*
- [3] Danny Cohen, "On Holy Wars and a Plea for Peace," IEN 137, 1 April 1980. *Already a classic document in data communications, this Internet Engineering Note gave the world the terms "little-endian" and "big-endian" with reference to byte- and bit-ordering.*
- [4] D. Waitzman, "A Standard for the Transmission of IP Datagrams on Avian Carriers," RFC 1149, 1 April 1990. *Yet Another IP encapsulation scheme, only this one gives new meaning to the idea of "best-effort" delivery.*
- [5] Poorer Richard and Professor Kynikos, "Gigabit Network Economics and Paradigm Shifts," RFC 1216, 1 April 1991. *Ultra-Low-Speed networking in its full glory.*

continued on next page

Internet Humor (*continued*)

- [6] Vinton Cerf, "Memo from the Consortium for Slow Commotion Research," RFC 1217, 1 April 1991. *Military applications of low-speed, low-efficiency networking systems.*
- [7] S. Greenfield, "Remembrance of Things Past," RFC 1300, February 1992. *A look-back at The Way Things Used To Be.*
- [8] Craig Partridge, "Today's Programming for KRFC AM 1313, Internet Talk Radio, RFC 1313," 1 April 1992. *Talk Radio with the leading lights of the Internet community. (See below).*

GARRETT A. WOLLMAN is a systems programmer, user consultant, student, and a lot of other things at the University of Vermont, Department of Computer Science and Electrical Engineering. He can be reached as: wollman@uvm.edu.

RFC 1313: Today's Programming for KRFC AM 1313 Internet Talk Radio

by Craig Partridge, BBN

Status of this memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

Welcome!

Hi and welcome to *KRFC Internet Talk Radio*, your place on the AM dial for lively talk and just-breaking news on internetworking. Sponsored by the Internet Society, KRFC serves the San Francisco Bay Area. For those of you outside the Bay Area, copies of program transcripts can be anonymously FTPed from `archives.krfc.com` the day after the program, or you can listen in via *vat*. Here's the programming for today, Wednesday, 1 April 1992:

Hacker's Hour with Phil Karn (Midnight): Phil's special guest today is Dr. David Mills, who will explain the special problems of correcting for the Doppler effect when trying to properly synchronize the new WWV receiver chip in your PC while flying on the Concorde.

Nighttime News (1 AM): Award winning Nighttime News gives you a full hour on those key facts you need to know before going to bed. Be sure to catch our network outage report with Elise Gerich. (Elise's report is sponsored by ANS).

Late At Night With Ole (2 AM): Call in your favorite Internetwork questions to Ole Jacobsen and his guests. Tonight's featured guests are John Moy, prime author of OSPF, and Milo Medin who will talk about how OSPF is great, but you really need to test it on 1822 networks to understand why.

Marty in the Morning (6 AM): Join the irrepressible Marty for five hours of eye-opening talk and commentary. Hear the latest on the commercial state of data networking in the US and who is at fault for limiting its growth. Special guest Kent England plans to drop by the studio today—listen in for the flames!

Education Report (11 AM): Gordon Cook solicits advice from Prof. David Farber on good ways to develop a research career. (In the likely event that Prof. Farber is unavailable at the last minute, Prof. Farber has arranged for Prof. David Sincoskie to take his place).

Lunch with Lynch (11:30 AM): Dan Lynch is on vacation this week and Vint Cerf is taking his place. Today Vint has lunch with Mitch Kapor of the EFF, MacArthur genius Richard Stallman, and Gen. Norman Schwarzkopf. Don't miss Vint's suggestions for wines to go with today's business lunch! [Lunch with Lynch is sponsored by Interop. Wines are provided by the vineyards in return for promotional considerations].

News (1 PM): Join Joyce and Jon as they report on the key networking news of the day. Don't miss their update on the latest address and port assignments and tips on upcoming RFCs!

Two by Four Time (2 PM): Today Marshall Rose will take out his two-by-four and apply it to Phill Gross for violating the Internet Standard Meeting Rules at the last IETF and starting a session before 9 AM. Additional victims to be announced. Today's show will be available as a book from Prentice Hall by next Tuesday.

Mike at the Mike (4 PM): Listen in to the Marina's favorite local DJ. Hear why They never listen and Never will! How come *The Book's* publishers don't seem to be able to add and why ATM is *Another Technical Mistake*. Then join MAP at 7:45 for a wee bit of this week's preferred single malt.

The Protocol Police (8 PM): Liven up your evening with the protocol police. Join our intrepid team of Stev Knowles and Mike St. Johns as they debug various TCP/IP implementations from the comfort of Mike's hot tub using Stev's water-proof portable PC. Last week they caught Peter Honeyman hijacking an NFS implementation. This week they're joined by Yakov Rekhter with his new Roto-Router tool, designed to catch routing anomalies. Who will our team nab this week?

Family Hour (10 PM): As part of this week's special series on children and networking, Bob Morris and Jerry Estrin talk about how much you should teach your young children about networking.

Securely Speaking (11 PM): Come eavesdrop as Steve Kent and Steve Crocker give you this weeks latest security news (if they're allowed to talk about it). And remember, just after 11 o'clock Steve and Steve will be reading this weeks encrypted message. If you're the first caller to call in with the right DES key to decrypt the message, you'll win \$1,000 and an all expenses paid trip to Ft. Meade! (US nationals only please).

Security considerations

Security issues are discussed in the above section.

Author's address

Craig Partridge
BBN
824 Kipling St
Palo Alto, CA 94301
Phone: 415-325-4541
E-Mail: craig@aland.bbn.com

Comments on “*Opinion: OSI Is (Still) a Good Idea*”

by Marshall T. Rose, Dover Beach Consulting, Inc.

Introduction

In the June *ConneXions*, Richard desJardins authored an opinion piece on “OSI is (Still) a Good Idea.” At “The Great OSI Debate” at INTEROP Spring 92, desJardins, Christian Huitema, and myself debated this topic. It was a target-rich environment and poor Richard was soundly trounced. Of course, this is hardly surprising since his job was to defend the indefensible, that is, he was to defend OSI.

Now, what will surprise you is that I actually agree with the title—but not the content—of desJardins’ opinion piece. You see, *Open Systems* really *are* a good idea. The problem is that OSI, as promulgated by the ISO and CCITT is a disaster. Before I respond to the points in his opinion piece (and explain, yet again, why poor Richard is in need of a reality check), let me explain myself.

Wrong approach

The concept of Open Systems is a good one. It is one in which we have a computer-communications market with products that are cost-effective and readily interoperate. No one argues this point. What we do argue is how to achieve this goal. desJardins believes that the ISO and CCITT are the appropriate venues. I believe this is plain wrong. The key thing to appreciate is that neither the ISO nor CCITT, nor their member bodies, view implementation, deployment, and interoperability as requirements for standardization. That is, they all develop standards on paper, which perhaps later are implemented after the “standard” imprint is issued. As we are fond of saying in the SNMP community:

The problems of the real world are remarkably resilient to administrative fiat.

What is meant by this is just because you write it down on paper doesn’t mean you can implement it well, or even implement it at all.

Dream or nightmare?

In contrast, the Internet community requires implementation, deployment, and interoperability to be demonstrated at each step of the standardization process. Although this process is hardly perfect, it does produce things called “standards” which do work. This is why my University mentor and colleague, Einar Stefferud says:

OSI is a beautiful dream, and TCP/IP is living it.

What is meant by this is that the process used by the Internet community is producing technology which is providing us with a competitive market for interoperable computer-communications products. This is a good thing, and I claim that this is derived from the way the Internet community produces standards. My corollary to Stef’s saying is:

OSI is a beautiful dream, and ISO/CCITT have turned it into a nightmare!

What is meant by this is that the “technology” produced by these organizations has not resulted in a competitive market. Rather it has resulted in a myriad of procrustean specifications which very rarely see the light of interoperable deployment.

Irrelevant

So, perhaps I am, as desJardins says, a bigot, by reminding people that the ISO/CCITT emperor has no clothes. If so, then I can only quote Mark Twain in that “Clothes make the man. Naked people have little or no effect on society.” And therein lies the delightful karma of the situation: the market has declared the ISO/CCITT approach to be irrelevant.

No one buys their version of OSI, because it simply doesn't work very well. So, to quote Stefferud one last time:

The bad news is that the ISO/CCITT have ignored the market, but the good news is that the market is ignoring them.

Predictions

And now, onto some modest desJardins-bashing! In his opinion piece, desJardins made seven "fearless" predictions. If the Editor of *ConneXions* would give me a few hundred column-inches, I'd be happy to explain why few of desJardins' predictions have even the remotest chance of coming true. Instead, I'm going to focus on just one. (Anyone with a keen sense of smell will be able to figure out why his other six predictions aren't that likely either. Here's what to do: for each one of his predictions see if you can find any empirical evidence to suggest a trend which would support a prediction. If so, then see if his rationale makes sense. If you try out this little intellectual exercise, I think you'll find yourself not very impressed with his "fearless" predictions.)

Skinny Stack hype

In terms of implementation, desJardins claims that "OSI will come free with UNIX by the end of the decade, and ... [that] Remote Procedure Calls [will] run like banshees over the OSI Skinny Stack." Let's ignore the fact that people (myself included [1]) have been predicting OSI bundled with UNIX for the last five years. Instead, let's concentrate on the latest bit of OSI marketing hype, the so-called *Skinny Stack*. The claim is that by implementing only the subset of the upper layer protocols that your application needs, your implementation will run faster. Somehow the few people actually foolhardy enough to try implementing the OSI upper layers are now vilified because they actually coded what the standards said. This is simply shameful. To begin with, Skinny Stack or not, it still takes five network transmissions before you can even send your first remote operation (a three-way handshake for transport and the session/presentation/application exchange), and then another four for tear-down. I don't care how thin you claim the code is going to be, the set-up and tear-down time is still going to kill performance. Second, anyone who thinks that a vendor is going to build a custom Skinny Stack for each application has obviously been spending too much time smoking standards, not implementing them.

Conclusion

So, enough desJardins-bashing for now. (That didn't hurt very much, did it, Richard?) In closing I want to do two things: first, I want to plug a paper:

M. T. Rose, "The Future of OSI: A Modest Prediction," Proceedings of the Conference on Upper Layer Protocols, Architectures, and Applications, pp. 356-365, IFIP Working Group 6.5, Vancouver, Canada, May 1992.

This explains my position in a lot more detail. The proceedings are published by North-Holland and should be available by early fall.

Second, I want to remind you that I agree with desJardins that OSI is a good idea. It always has been. Where we disagree is the means whereby this goal can be achieved. I claim that only *de facto* standards have any value; and, that the only way to produce those is through a cycle in which implementation, deployment, and interoperability are demonstrated throughout. desJardins believes in *de jure* standards as promulgated by the ISO and CCITT. I leave it to you to decide whether you prefer the dream or the nightmare!

References

- [1] Rose, M. T., "OSI protocols within an openly available, POSIX-conformant, Berkeley UNIX environment," *ConneXions*, Volume 2, No. 10, October 1988.

Book Review

Zen and the Art of the Internet: A Beginner's Guide to the Internet, (electronic document).

Goal *Zen and the Art of the Internet* is a new guide to the Internet that was written by Brendan Kehoe of Widener University. His goal was to introduce the reader to the resources that are available on the Internet. At the same time, Kehoe tried to avoid system specific information. It should be noted that parts of *Zen and the Art of the Internet* were derived from other works.

Zen and the Art of the Internet starts off with a chapter on network basics. This chapter is a good introduction to the Internet, but it is not a general guide to networking. Rather, it is Internet and TCP/IP specific. If this chapter can be faulted for anything, it is that it oversimplifies some of the material. On the other hand, it definitely should not scare off the novice user.

E-mail and FTP The e-mail and FTP chapters are very good, although they do get technical at times. The e-mail chapter could be improved by the addition of a section on etiquette similar to the excellent one in the FTP chapter.

Telnet The *Telnet* chapter is packed with examples of *Telnet*-accessible services, and it explains how to find out about more services. I was rather disappointed by the omission of any information on *tn3270*. A description of how *Telnet* is different on IBM mainframes is also needed. These omissions may lead to some confusion on the part of IBM mainframe users.

Kehoe describes other tools that are available on the Internet. These descriptions are well-rounded and useful, but Kehoe has just covered the most common tools.

FAQ One of the most outstanding sections of *Zen and the Art of the Internet* is called "Things You'll Hear About." In a lot of ways, this chapter is a FAQ (*Frequently Asked Questions*) to the Internet, and it will answer many questions of the new network user. At the same time, it introduces the novice user to the folklore of the Internet without being intimidating.

Zen and the Art of the Internet also has useful sections that contain information about commercial services, other networks, how to retrieve files, and how to find out more about the Internet. The USENET chapter does a great job of covering the most common misconceptions people have about that network. The document includes a helpful glossary.

The conclusion states "this guide is far from complete—the Internet changes on a daily (if not hourly) basis." Then Kehoe goes on to ask for suggestions. For *Zen and the Art of the Internet* to be useful in the long run, it will need to be updated on a fairly regular basis. From what I can tell, it sounds like Kehoe is planning on doing this. I'm sending in my suggestions, and I highly recommend you do the same.

Impressive Overall, I was very impressed with this document. In fact, the same day that I downloaded it I had our receptionist make copies and distribute them to the whole Academic Computing Support Staff. In a couple of days, I am going to do the same with our library. My girlfriend's university just got on the Internet and I'm giving her two sources of information to start with: the first is HYTELNET and the second is going to be *Zen and the Art of the Internet*.

It has a few rough spots, but I'm sure that Kehoe will fix them. The biggest problem is that it paints too rosy a picture of the Internet, but this kind of document is intended to get users interested in the network not to critique it.

I try to stay ahead of most Internet users in terms of my knowledge of what's available and how to access it. Well, I learned a couple of things while reading *Zen and the Art of the Internet*, so it is not just for novices. At the same time, it is easily understandable by novices. My message to Brendan Kehoe is: "Keep up the good work!"

Access Instructions

The file is available on host FTP.CS.WIDENER.EDU(147.31.254.132) in the directory pub/zen and on FTP.UU.NET (137.39.1.9) in the directory /inet/doc. Although the author reports that he has signed an agreement with a major publishing house, he has indicated that the network versions will continue to be available.

Acknowledgments

Roy Tennant, the University of California Berkeley, and Charles Bailey, University of Houston, for Access Instructions and original editing of this review.

—Billy Barron, University of North Texas. © 1992

[This review is reprinted with permission from *The Public-Access Computer Systems Review*, Volume 3, No. 1].

NorthWestNet User Services Internet Resource Guide

In the fall of 1991, NorthWestNet (an NSFNET regional network) published the third edition of its *NorthWestNet User Services Internet Resource Guide* (NUSIRG). Response to NUSIRG has far exceeded NorthWestNet's expectations. The manual is now in its fourth printing, and nearly 1000 copies have been distributed. One regional network recently placed a bulk order for 200 copies for its members.

Contents

NUSIRG serves as a guide to such Internet services and resources as electronic mail, file transfer (FTP), remote login (Telnet), discussion groups, on-line library catalogues, and supercomputer access. A typical chapter includes the definition of a resource, examples and a step-by-step explanation of how it is used, and suggestions on where to find more information. NUSIRG also contains lists and tables of information such as LISTSERV discussion groups, FTP sites, electronic mail gateways, supercomputer sites, and multi-disciplinary databases.

Tutorial and reference

This 300-page manual is written both as a how-to tutorial for the Internet beginner and as a reference manual for the more experienced user. As NUSIRG author Jonathan Kochmer says in the manual, he has tried to be "thorough without being thoroughly overwhelming."

NorthWestNet plans to continue updating NUSIRG to keep up with the rapid evolution of Internet resources and applications. In fact, research has already begun on the fourth edition.

Ordering information

For a bound copy of NUSIRG, send \$20.00 (plus 8.2% sales tax for orders in Washington State) to:

NorthWestNet
15400 SE 30th Place, Suite 202,
Bellevue, WA 98007.

Please make checks payable to NorthWestNet.

Conference Announcement

NSC'92—The Network Services Conference 1992 will be held in Pisa, Italy, November 3–5, 1992. The conference is being organized by EARN in conjunction with EUnet/EurOpen, NORDUnet, RARE, and RIPE.

Overview

The world of academic and research networking has evolved to the point where the protocol wars have become largely irrelevant. This is demonstrated by the recent appearance of high-level networking tools which are worldwide in scope and which run simultaneously over many different lower layers.

NSC 92 will focus on issues in providing services to customers, with special attention paid to the recent and exciting developments in new global high-level tools such as *World-Wide Web*, *Prospero*, *Archie*, *Alex*, *Gopher*, and *WAIS*. We will address the impact of the new global tools on service development and support, the changing function of traditional tools and services (such as archives), upcoming specific services such as new databases, and the future role of the library. User support at the campus level, and the role of support in accessing global services, will be addressed.

The conference will be of greatest interest to network service providers and sophisticated users who are changing their focus from providing or obtaining bandwidth to offering, supporting, and using varied and powerful services. Talks and other conference activities will address the needs of the research, academic, educational, governmental, industrial, and commercial network communities.

Venue

Pisa is situated in Tuscany on the Arno river. The Italian poet Gabriele D'Annunzio named Pisa's Piazza della Torre: "The Square of Miracles," and yet the definition could be extended with equal justice to the whole city. Pisa is not only an art center with few rivals; it is steeped in culture and science and offers an up-to-date infrastructure. The conference will be held at the Palazzo dei Congressi, near the city center and at walking distance to the Hotels.

Topics

Papers for presentation at the conference have been solicited in the following areas (the submission deadline was May 31, 1992):

- *Dealing with the Information Explosion:*
 - New Global Information Access Tools
 - Utilizing Established Information Access and Distribution Tools
- *Managing Global Network Information Services:*
 - Coordination/Duplication, Security, Privacy, Authentication
 - Closed Group Applications
- *The Electronic Library:*
 - Local Databases, Remote Databases, OPACS, CD-ROMS
 - Inter-Library Cooperation
- *User/Customer Support:*
 - Help Desks, Documentation, Reaching the Customers
- *Assessing Customer Needs*
- *Special Interest Communities*

- *Group Communication Technologies & Services—“Groupware”*
- *Networking for Schools*
- *Delivering Messaging to the Desktop:*
 - Practical Experiences, Products, Security, Interface issues
- *Beyond ASCII:*
 - Character sets, Multimedia
 - Creating, Encoding, Receiving
- *Economic Aspects of Networking:*
 - Bandwidth, E-mail Access, Efficiency, Control
 - Recent European Networking Developments

Posters and demonstrations

A poster wall will be available to participants for the display of their posters and projects. A terminal room with connectivity to EARN and the Internet will be available to delegates. A room will be available for workstations and PCs to be used for demonstrations. An Ethernet connected to the Internet will be available in the room. Connectivity to the Internet will be via a 64Kbps line to CNUCE. The minimum bandwidth between CNUCE and CERN is 512Kbps. People interested in setting up demonstrations may send their questions to:

NSCINFO@FRORS12.BITNET

Further information

The conference program with information on how to register will be distributed with the second announcement around 1 August 1992. Further information will be available through an ad hoc conference mailing list. If you want to make sure you receive the invitation, as well as the preliminary program, please ask for subscription to the conference mailing list (NSC92@FRORS12.BITNET) sending mail, e-mail or fax specifying your e-mail address to:

Nadine Grange
EARN Office
c/o CIRCE
BP 167
F91403 Orsay
FRANCE
Tel: +33 1 6982 3973
Fax: +33 1 6928 5273
E-mail: GRANGE@FRORS12.BITNET

(General inquiries can be made at NSCINFO@FRORS12.BITNET)

Program Committee

Dennis Jennings, Ireland (Chair); Rob Blokzijl, The Netherlands; Daniele Bovio, France; Paul Bryant, United Kingdom; Avi Cohen, Israel; Laszlo Csaba, Hungary; Hans Deckers, France; Jean-Loic Delhayé, France; Jill Foster, United Kingdom; Frode Greisen, Denmark; Glenn Kowack, the Netherlands; Stelios Orphanoudakis, Greece; David Sitman, Israel; Stefano Trumpy, Italy.

Organizing Committee

Frode Greisen, Denmark (Chair); Hans Deckers, France; Dennis Jennings, Ireland; Glenn Kowack, the Netherlands; Marco Sommani, Italy.

Internet Society Takes Historic Steps

At its first formal meeting at INET '92 in June, the Internet Society Board of Trustees gathering from all corners of the world in Kobe, Japan, took several historic actions that will significantly affect inter-networking worldwide.

The Internet Society (ISOC) is an international professional organization established for evolving and extending availability of the techniques and technologies that allow diverse information systems to openly communicate. It also includes the huge network of networks known as the Internet which links millions of users worldwide. These technologies and the Internet are very rapidly evolving and widely viewed as critically important infrastructure.

New name for the IAB

The organization formerly known as the Internet Activities Board was merged into the Internet Society as a body called the *Internet Architecture Board* (IAB). The IAB evolves the technology of the Internet and develops the series of international standards which are the predominant means today for common open interconnection, management, and use of diverse equipment, networks and applications. Some of the most popular include electronic mail, file transfer, news distribution, remote logon, and knowledge discovery.

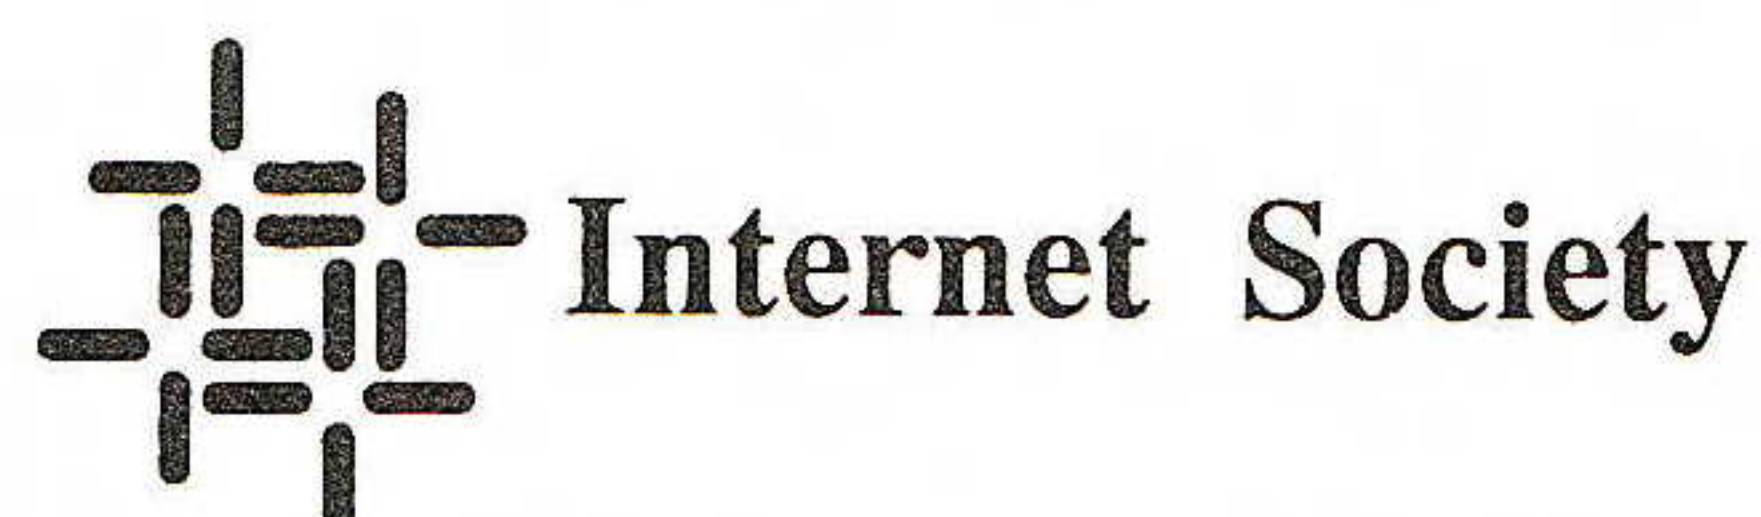
Collaboration

The Board also decided to establish a cooperative relationship with the Geneva-based U.N. specialized agency known as the *International Telecommunication Union* (ITU). In conjunction with this action, the Board also decided to establish an advisory liaison relationship between the standards bodies of the two organizations, thus effecting a link between the IAB and the International Telegraph and Telephone Consultative Committee (CCITT). Lastly the ISOC will submit a contribution to the ITU's Plenipotentiary Conference that underscores internetworking as critical infrastructure and the use of the Internet to significantly enhance global telecommunications collaboration.

By taking these historic steps, the Board hopes not only to greatly advance the development and use of internetworking technologies and the Internet, but also to stimulate more effective collaboration within the telecommunications field.

For further information, contact:

Internet Society Secretariat
1895 Preston White Drive
Suite 100
Reston, VA 22091
USA
Tel: +1 703 620 8990
Fax: +1 703 620 0913
E-mail: isoc@nri.reston.va.us



Ed.: We will have a more detailed report from INET '92 in a future issue.

Upcoming Events

The *First Symposium on High Performance Distributed Computing* (HPDC-1) will be held September 9–10, 1992, at Sheraton University, Syracuse, New York. The symposium is sponsored by the IEEE Computer Society, The New York State Center for Advanced Technology in Computer Applications and Software Engineering at Syracuse University, and Northeast Parallel Architectures Center at Syracuse University. Symposium proceedings will be published by the IEEE Computer Society.

Theme The theme of this workshop is to investigate software techniques and architectural support for application of parallel and distributed computing for solving computationally intensive applications across a network of high-performance computers.

Topics Papers that address all aspects of achieving parallel and distributed computing across a high speed network of computers have been sought. (The deadline for submission was May 1, 1992). Papers that deal with experimental results and prototypes were expected in the following areas:

- Parallel and distributed algorithms to solve computationally intensive problems across a LAN, MAN, or WAN
- Architectural support for high-speed communications or interconnection networks
- Gigabit network architectures
- Networking for multimedia data
- High-speed communication transport protocols to achieve Gigabit per second application-to-application transfer rates
- Performance evaluation of experimental systems to solve supercomputing applications across a network of computers

More information For more information contact the Workshop General Chair:

Geoffrey Fox
NPAC
Syracuse University
315-443-4741 • hpdc@nova.npac.syr.edu

Symposium proceedings will be published by the IEEE Computer Society.

Write to *ConneXions*!

Have a question about your subscription? Are you moving, and need to give us your new address? Suggestions for topics? Want to write an article? An *Opinion* piece? A letter to the Editor? Have a question for an author? Need a *ConneXions* binder? Want to enquire about back issues? (there are now over sixty-four to choose from; ask for our free 1987–1992 index booklet). We want to hear from you. Simply write, call, fax, or e-mail us at:

ConneXions—The Interoperability Report
480 San Antonio Road, Suite 100
Mountain View, CA 94040–1219
USA

Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)

Fax: +1 415-949-1779

E-mail: connexions@interop.com

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President,
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute

Subscribe to CONNEXIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp.Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

Back issues available upon request \$15./each
Volume discounts available upon request

480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779
connexions@interop.com

CONNEXIONS